

Runtime enforcement of reactive systems using synchronous enforcers^{*}

Srinivas Pinisetty¹, Partha S Roop², Steven Smyth³, Stavros Tripakis^{1,4},
Reinhard von Hanxleden³

¹ Aalto University, Finland `First.Last@aalto.fi`

² University of Auckland, New Zealand `p.roop@aucklanduni.ac.nz`

³ University of Kiel, Germany `ssm,rvh@informatik.uni-kiel.de`

⁴ University of California, Berkeley, USA

Abstract. Synchronous programming is a paradigm of choice for the design of safety-critical reactive systems. Runtime enforcement is a technique to ensure that the output of a black-box system satisfies some desired properties. This paper deals with the problem of runtime enforcement in the context of synchronous programs. We propose a framework where an enforcer monitors both the inputs and the outputs of a synchronous program and (minimally) edits erroneous inputs/outputs in order to guarantee that a given property holds. We define enforceability conditions, develop an online enforcement algorithm, and prove its correctness. We also report on an implementation of the algorithm on top of the KIELER framework for the SCCharts synchronous language. Experimental results show that enforcement has minimal execution time overhead, which decreases proportionally with larger benchmarks.

1 Introduction

Runtime verification (RV) [11,7] is an active area of research on methods that dynamically verify a set of desirable properties over an execution of a “black-box” system. An alternative to such passive runtime analysis is runtime enforcement (RE) [19,8,12,15]. In RE mechanisms, an *enforcer* is synthesized to observe the executions of a black-box system to ensure that a set of desired properties are satisfied. In the event of a violation, the enforcer performs certain evasive actions so as to prevent the violation. The evasive actions might include blocking the execution [19], modifying input sequence by suppressing and / or inserting actions [12], and buffering input actions until a future time when it could be forwarded [8,15]. These enforcement mechanisms are not suitable for synchronous reactive systems since delaying the reaction or terminating the system is infeasible. Considering this, there is recent interest in runtime enforcement of synchronous reactive systems [5].

A synchronous reactive system is non-terminating and interacts continuously with the adjoining environment. Hence, the system execution may be considered as a series of steps, where in each step the system reads the inputs from the environment, calls a *reaction function* that computes the outputs for emission. Synchronous programming languages [3] are well suited for the design of synchronous reactive systems. They use observers [9] to express safety properties, which are verified statically (using model checking). There have also been limited attempts to use observers as runtime entities [17], for example for automatic test case generation. More recently Rushby studies applications of observers for the expression of assumptions and axioms in addition to test case generation [18]. However, there have been no studies on the bi-directional RE problem for synchronous reactive systems, which is the focus of the current paper.

We consider bi-directional RE of synchronous programs, and the general context is illustrated in Figure 1. Here, $\{i_1, i_2, \dots, i_n\}$ are inputs from the environment to the enforcer, $\{i'_1, i'_2, \dots, i'_n\}$ are transformed inputs from the enforcer to the program, $\{o_1, o_2, \dots, o_m\}$ are outputs of the program to the enforcer, and $\{o'_1, o'_2, \dots, o'_m\}$ are transformed outputs from the enforcer to the environment. RE for synchronous reactive systems is distinct from the existing RE mechanisms such as [8,12,15,19] since the enforcement mechanism for a synchronous reactive system cannot halt the system or delay events, and must react instantaneously when an error is observed. Moreover, we consider bi-directional enforcement where the enforcer needs to consider the status of the environment and the program in order to enforce the policies. The enforcer must respect the *causality* aspects i.e. every reactive cycle must start with the

^{*} This work has been partially supported by the Academy of Finland, the U.S. National Science Foundation (awards #1329759 and #1139138), and the Deutsche Forschungsgemeinschaft (PRETSY2 project, award DFG HA 4407/6-2).




Fig. 1: Bi-directional enforcement for synchronous programs.

environment, where the status of the environment inputs must determine the reaction. After the program has reacted, the generated outputs are emitted to the environment. Considering this, the enforcer must act as an intermediary such that it first intercepts the inputs from the environment to validate them relative to the policy and forward the inputs to the program once the policy is satisfied. In the event of any violation, the enforcer may suitably alter the inputs before forwarding to the program. After the program has reacted to these inputs, again the enforcer must ensure that either the policy is satisfied and hence the outputs are forwarded unchanged to the environment or a violation has happened that needs to be handled by altering the outputs to prevent policy violation.

We study the problem of synthesizing an enforcer for any given safety property φ . Similar to enforcement mechanisms in [19,8,12,15], several constraints are required on how an enforcer transforms input-output words. The enforcer cannot delay events, and cannot block execution, but it is allowed to *edit* an event when necessary (i.e., when the event that it receives as input leads to a violation). The notions of *soundness* and *transparency* are similar to the existing enforcement mechanisms [19,8,12,15], where soundness means that the output of the enforcer must satisfy property φ , and transparency expresses that the enforcer should not modify events unnecessarily. In the proposed framework, we also introduce additional requirements called *causality*, and *instantaneity*. These constraints are developed specifically to respect synchronous execution, detailed in Section 3.

Contributions. In this paper, we study and formally define, for the first time, the bi-directional enforcer synthesis problem for synchronous reactive systems (expressed as synchronous programs). The main contributions of the paper are (1) We formally define the bi-directional enforcer synthesis problem and characterize the set of safety properties which can be enforced (Section 3), (2) We develop an enforcement algorithm (Section 4) and prove its correctness, (3) We report on an implementation of the algorithm on top of the KIELER framework for the SCCharts synchronous language (Section 5), and (5) We evaluate the approach over a range of synchronous programs in the SCCharts language [21] to illustrate scalability and practicality (Section 5).

2 Preliminaries and Notation

A finite (resp. infinite) word over a finite alphabet Σ is a finite sequence $\sigma = a_1 \cdot a_2 \cdots a_n$ (resp. infinite sequence $\sigma = a_1 \cdot a_2 \cdots$) of elements of Σ . The set of finite (resp. infinite) words over Σ is denoted by Σ^* (resp. Σ^ω). The *length* of a finite word σ is n and is noted $|\sigma|$. The empty word over Σ is denoted by ϵ_Σ , or ϵ when clear from the context. The *concatenation* of two words σ and σ' is denoted as $\sigma \cdot \sigma'$. A word σ' is a *prefix* of a word σ , denoted as $\sigma' \preceq \sigma$, whenever there exists a word σ'' such that $\sigma = \sigma' \cdot \sigma''$; σ is said to be an *extension* of σ' .

We consider a reactive system with a finite ordered sets of Boolean inputs $I = \{i_1, i_2, \dots, i_n\}$ and Boolean outputs $O = \{o_1, o_2, \dots, o_m\}$. The input alphabet is $\Sigma_I = 2^I$, and the output alphabet is $\Sigma_O = 2^O$ and the input-output alphabet $\Sigma = \Sigma_I \times \Sigma_O$. Each input (resp. output) event will be denoted as a bit-vector/complete monomial. For example, let $I = \{A, B\}$. Then, the input $\{A\} \in \Sigma_I$ is denoted as 10, while $\{B\} \in \Sigma_I$ is denoted as 01 and $\{A, B\} \in \Sigma_I$ is denoted as 11. A reaction (or input-output event) is of the form (x_i, y_i) , where $x_i \in \Sigma_I$ and $y_i \in \Sigma_O$.

Given an input-output word $\sigma = (x_1, y_1) \cdot (x_2, y_2) \cdots (x_n, y_n) \in \Sigma^*$, the input word obtained from σ is $\sigma_I = x_1 \cdot x_2 \cdots x_n \in \Sigma_I$ which is the projection on inputs ignoring outputs. Similarly, the output word obtained from σ is $\sigma_O = y_1 \cdot y_2 \cdots y_n \in \Sigma_O$ is the projection on outputs.

An execution σ of a synchronous program \mathcal{P} is an infinite sequence of input-output events $\sigma \in \Sigma^\omega$, and the *behavior* of a synchronous program \mathcal{P} is denoted as $exec(\mathcal{P}) \subseteq \Sigma^\omega$. The *language* of \mathcal{P} is denoted by $\mathcal{L}(\mathcal{P}) = \{\sigma \in \Sigma^* \mid \exists \sigma' \in exec(\mathcal{P}) \wedge \sigma \preceq \sigma'\}$ i.e. $\mathcal{L}(\mathcal{P})$ is the set of all finite prefixes of the sequences in $exec(\mathcal{P})$.

A property φ over Σ defines a set $\mathcal{L}(\varphi) \subseteq \Sigma^*$. A program $\mathcal{P} \models \varphi$ iff $\mathcal{L}(\mathcal{P}) \subseteq \mathcal{L}(\varphi)$. Given a word $\sigma \in \Sigma^*$, $\sigma \models \varphi$ iff $\sigma \in \mathcal{L}(\varphi)$. A property φ is *prefix-closed* if all prefixes of all words from $\mathcal{L}(\varphi)$ are also in $\mathcal{L}(\varphi)$: $\mathcal{L}(\varphi) = \{w \mid \exists w' \in \mathcal{L}(\varphi) : w \preceq w'\}$. In this paper, we consider prefix-closed properties. Properties are formally expressed as safety automata that we define in the sequel.

Definition 1 (Safety Automaton). A safety automaton (SA) $\mathcal{A} = (Q, q_0, q_v, \Sigma, \rightarrow)$ is a tuple, where Q is the set of states, called *locations*, $q_0 \in Q$ is an unique initial location, $q_v \in Q$ is a unique violating (non-accepting) location, $\Sigma = \Sigma_I \times \Sigma_O$ is the alphabet, and $\rightarrow \subseteq Q \times \Sigma \times Q$ is the transition relation. All the locations in Q except q_v (i.e., $Q \setminus \{q_v\}$) are accepting locations. Location q_v is a unique non-accepting (trap) location, and there are no transitions in \rightarrow from q_v to a location in $Q \setminus \{q_v\}$. Whenever there exists $(q, a, q') \in \rightarrow$, we denote it as $q \xrightarrow{a} q'$. Relation \rightarrow is extended to words $\sigma \in \Sigma^*$ by noting $q \xrightarrow{\sigma, a} q'$ whenever there exists q'' such that $q \xrightarrow{\sigma} q''$ and $q'' \xrightarrow{a} q'$. A location $q \in Q$ is *reachable* from q_0 if there exists a word $\sigma \in \Sigma^*$ such that $q_0 \xrightarrow{\sigma} q$.

An SA $\mathcal{A} = (Q, q_0, q_v, \Sigma, \rightarrow)$ is *deterministic* if $\forall q \in Q, \forall a \in \Sigma, (q \xrightarrow{a} q' \wedge q \xrightarrow{a} q'') \implies (q' = q'')$. \mathcal{A} is *complete* if $\forall q \in Q, \forall a \in \Sigma, \exists q' \in Q, q \xrightarrow{a} q'$. A word σ is *accepted* by \mathcal{A} if there exists $q \in Q \setminus \{q_v\}$ such that $q_0 \xrightarrow{\sigma} q$. The set of all words accepted by \mathcal{A} is denoted as $\mathcal{L}(\mathcal{A})$.

Remark 1. In the rest of this paper, φ is a safety property defined as deterministic and complete SA $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$. If the user provides an non-deterministic or incomplete automaton, we determinize and complete it first. We also consider that Q does not contain any (redundant) locations that are unreachable from q_0 .

Due to the causality requirement, the enforcer has to first transform inputs from the environment in each step according to property φ defined as SA \mathcal{A}_φ . We thus need to consider the input property that we obtain from \mathcal{A}_φ by projecting on inputs.

Definition 2 (Input safety automaton \mathcal{A}_{φ_I}). Given $\varphi \subseteq \Sigma^*$, defined as SA $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$, input SA $\mathcal{A}_{\varphi_I} = (Q, q_0, q_v, \Sigma_I, \rightarrow_I)$ is obtained from \mathcal{A}_φ by ignoring outputs on the transitions, i.e., for every transition $q \xrightarrow{(x,y)} q' \in \rightarrow$ where $(x, y) \in \Sigma$, there is a transition $q \xrightarrow{x} q' \in \rightarrow_I$, where $x \in \Sigma_I$. $\mathcal{L}(\mathcal{A}_{\varphi_I})$ is denoted as $\varphi_I \subseteq \Sigma_I^*$.




Fig. 2: SA (left), and its input SA (right).

Example 1 (Example property defined as SA and its input SA). Let $I = \{A, B\}$ and $O = \{R\}$. Consider the following property: S_1 : “A and B cannot happen simultaneously, and also B and R cannot happen simultaneously”. The safety automaton in Figure 2a defines property S_1 . Figure 2b presents the input SA for the SA in Figure 2a defining property S_1 . Though the SA \mathcal{A}_φ is deterministic, the input SA \mathcal{A}_{φ_I} might be non-deterministic as is the case in Figure 2b.

Lemma 1. Let $\mathcal{A}_{\varphi_I} = (Q, q_0, q_v, \Sigma_I, \rightarrow_I)$ be the input automaton obtained from $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$. We have the following properties:

- 1 $\forall (x, y) \in \Sigma, \forall q, q' \in Q : q \xrightarrow{(x,y)} q' \implies q \xrightarrow{x}_I q'$.
- 2 $\forall x \in \Sigma_I, \forall q, q' \in Q : q \xrightarrow{x}_I q' \implies \exists y \in \Sigma_O : q \xrightarrow{(x,y)} q'$.

Intuitively, property 1 of Lemma 1 states that if there is a transition from state $q \in Q$ to state $q' \in Q$ upon input-output event $(x, y) \in \Sigma$ in the automaton \mathcal{A}_φ , then there is also a transition from state q to state q' in the input automaton \mathcal{A}_{φ_I} upon the input event $x \in \Sigma_I$. Property 2 of Lemma 1 states that if there is a transition from state $q \in Q$ to state $q' \in Q$ upon input event $x \in \Sigma_I$, then there certainly exists an output event $y \in \Sigma_O$ s.t. there is a transition from state q to state q' upon event (x, y) in the automaton \mathcal{A}_φ . Lemma 1 immediately follows from Definitions 1 and 2.

Edit Functions Consider property $\varphi \subseteq \Sigma^*$ defined as SA $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$, and its input SA $\mathcal{A}_{\varphi_I} = (Q, q_0, q_v, \Sigma_I, \rightarrow_I)$ obtained from \mathcal{A}_φ by projecting on inputs. We introduce editl_{φ_I} (resp. editO_φ), that the enforcer uses for editing input (resp. output) events (when necessary), according to input property φ_I (resp. property φ).

- **$\text{editl}_{\varphi_I}(\sigma_I)$** : Given $\sigma_I \in \Sigma_I^*$, $\text{editl}_{\varphi_I}(\sigma_I)$ is the set of input events x in Σ_I such that the word obtained by extending σ_I with x satisfies property φ_I . Formally,

$$\text{editl}_{\varphi_I}(\sigma_I) = \{x \in \Sigma_I : \sigma_I \cdot x \models \varphi_I\}.$$

Considering the SA $\mathcal{A}_{\varphi_I} = (Q, q_0, q_v, \Sigma_I, \rightarrow_I)$, the set of events in Σ_I that allow to reach a state in $Q \setminus \{q_v\}$ from a state $q \in Q \setminus \{q_v\}$ is defined as:

$$\text{editl}_{\mathcal{A}_{\varphi_I}}(q) = \{x \in \Sigma_I : q \xrightarrow{x}_I q' \wedge q' \neq q_v\}.$$

For example, consider the SA in Figure 2b obtained from the SA in Figure 2a by ignoring outputs. Let $\sigma = (10, 0) \cdot (01, 1)$, and thus $\sigma_I = 10 \cdot 01$. Then, $\text{editl}_{\varphi_I}(\sigma_I) = \Sigma_I \setminus \{11\}$. Also, $q_0 \xrightarrow{10 \cdot 01}_I q_0$, and $\text{editl}_{\mathcal{A}_{\varphi_I}}(q_0) = \Sigma_I \setminus \{11\}$.

- **$\text{nondet-editl}_{\mathcal{A}_{\varphi_I}}(q)$** : If $\text{editl}_{\mathcal{A}_{\varphi_I}}(q)$ is non-empty, then $\text{nondet-editl}_{\mathcal{A}_{\varphi_I}}(q)$ returns an element (chosen randomly) from $\text{editl}_{\mathcal{A}_{\varphi_I}}(q)$, and is undefined if $\text{editl}_{\mathcal{A}_{\varphi_I}}(q)$ is empty.
- **$\text{editO}_\varphi(\sigma, x)$** : Given an input-output word $\sigma \in \Sigma^*$ and an input event $x \in \Sigma_I$, $\text{editO}_\varphi(\sigma, x)$ is the set of output events y in Σ_O s.t. the input-output word obtained by extending σ with (x, y) satisfies property φ . Formally,

$$\text{editO}_\varphi(\sigma, x) = \{y \in \Sigma_O : \sigma \cdot (x, y) \models \varphi\}.$$

Considering the automaton $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$ defining property φ , and an input event $x \in \Sigma_I$, the set of output events y in Σ_O that allow to reach a state in $Q \setminus \{q_v\}$ from a state $q \in Q \setminus \{q_v\}$ with (x, y) is defined as:

$$\text{editO}_{\mathcal{A}_\varphi}(q, x) = \{y \in \Sigma_O : q \xrightarrow{(x, y)} q' \wedge q' \neq q_v\}.$$

For example, consider property S_1 defined by the automaton in Figure 2a. We have $\text{editO}_{\mathcal{A}_\varphi}(q_0, 01) = \{0\}$.

- **$\text{nondet-editO}_{\mathcal{A}_\varphi}(q, x)$** : If $\text{editO}_{\mathcal{A}_\varphi}(q, x)$ is non-empty, then $\text{nondet-editO}_{\mathcal{A}_\varphi}(q, x)$ returns an element (chosen randomly) from $\text{editO}_{\mathcal{A}_\varphi}(q, x)$, and is undefined if $\text{editO}_{\mathcal{A}_\varphi}(q, x)$ is empty.

3 Problem Definition

In this section, we formalize the RE problem for synchronous programs. In the setting we consider, as illustrated in Figure 1, an enforcer monitors and corrects both inputs and outputs of a synchronous program according to a given safety property $\varphi \subseteq \Sigma^*$. We assume that the “black-box” synchronous program may be invoked through a special function call called `ptick`, which is invoked exactly once during each reaction / synchronous step. Formally, `ptick` is a function from Σ_I to Σ_O that takes a bit vector $x \in \Sigma_I$ and returns a bit vector $y \in \Sigma_O$.

An enforcer for a property φ can only edit an input-output event when necessary, and it cannot block, delay or suppress events. Let us recall the two functions editl_{φ_I} and editO_φ that were introduced in Section 2 that the enforcer for φ uses to edit the current input (respectively output) event according to the property φ . At an abstract level, an enforcer can be seen as a function that transforms input-output words. An enforcement function for a given property φ takes as input an input-output word over Σ and outputs an input-output word over Σ that belongs to φ .

Definition 3 (Enforcer for φ). Given property $\varphi \subseteq \Sigma^*$, an enforcer for φ is a function $E_\varphi : \Sigma^* \rightarrow \Sigma^*$ satisfying the following constraints:

Soundness

$$\forall \sigma \in \Sigma^* : E_\varphi(\sigma) \models \varphi. \quad (\text{Snd})$$

Monotonicity

$$\forall \sigma, \sigma' \in \Sigma^* : \sigma \preceq \sigma' \Rightarrow E_\varphi(\sigma) \preceq E_\varphi(\sigma'). \quad (\text{Mono})$$

Instantaneity

$$\forall \sigma \in \Sigma^* : |\sigma| = |E_\varphi(\sigma)|. \quad (\text{Inst})$$

Transparency

$$\begin{aligned} \forall \sigma \in \Sigma^*, \forall x \in \Sigma_I, \forall y \in \Sigma_O : \\ E_\varphi(\sigma) \cdot (x, y) \models \varphi \implies E_\varphi(\sigma \cdot (x, y)) = E_\varphi(\sigma) \cdot (x, y). \end{aligned} \quad (\text{Tr})$$

Causality

$$\begin{aligned} \forall \sigma \in \Sigma^*, \forall x \in \Sigma_I, \forall y \in \Sigma_O, \exists x' \in \text{editl}_{\varphi_1}(E_\varphi(\sigma)_I), \\ \exists y' \in \text{editO}_\varphi(E_\varphi(\sigma), x') : E_\varphi(\sigma \cdot (x, y)) = E_\varphi(\sigma) \cdot (x', y'). \end{aligned} \quad (\text{Cau})$$

The input-output sequence released as output by the enforcer upon reading the input-output sequence σ is $E_\varphi(\sigma)$, and $E_\varphi(\sigma)_I \in \Sigma_I^*$ is the projection on the inputs. Note, $\text{editl}_{\varphi_1}(E_\varphi(\sigma)_I)$ returns a set of input events in Σ_I , s.t. $E_\varphi(\sigma)_I$ (which is the projection of input-output word $E_\varphi(\sigma)$ to the input alphabet) followed by any event in $\text{editl}_{\varphi_1}(E_\varphi(\sigma)_I)$ satisfies φ_I . $\text{editO}_\varphi(E_\varphi(\sigma), x')$ returns a set of output events in Σ_O , s.t. for any event y in $\text{editO}_\varphi(E_\varphi(\sigma), x')$, $E_\varphi(\sigma) \cdot (x', y)$ satisfies φ .

- Soundness (**Snd**) means that for any word $\sigma \in \Sigma^*$, the output of the enforcer $E_\varphi(\sigma)$ must satisfy φ .
- Monotonicity (**Mono**) expresses that the output of the enforcer for an extended word σ' of a word σ , extends the output produced by the enforcer for σ . The monotonicity constraint means that the enforcer cannot undo what is already released as output.
- Instantaneity (**Inst**) expresses that for any given input-output word σ as input to the enforcer, the output of the enforcer $E_\varphi(\sigma)$ should contain exactly the same number of events that are in σ (i.e., E_φ is length-preserving). This means that the enforcer cannot delay, insert and suppress events. Whenever the enforcer receives a new event, it has to react instantaneously and has to produce an output event immediately.
- Transparency (**Tr**) expresses that for any given word σ and any event (x, y) , if the output of the enforcer for σ (i.e., $E_\varphi(\sigma)$) followed by the event (x, y) satisfies the property φ (i.e., $E_\varphi(\sigma) \cdot (x, y) \models \varphi$), then the output that the enforcer produces for input $\sigma \cdot (x, y)$ will be $E_\varphi(\sigma) \cdot (x, y)$. This means that the enforcer makes no change when no change is needed in order to satisfy the property φ .
- Causality (**Cau**) expresses that for every input-output event (x, y) the enforcer produces input-output event (x', y') where the enforcer first processes the input part x , to produce the transformed input x' according to property φ using editl_{φ_1} . The enforcer later reads and transforms output $y \in \Sigma_O$ which is the output of the program after invoking function ptick with the transformed input x' , to produce the transformed output y' using editO_φ .

Remark 2. Let $E_\varphi(\sigma)$ be the input-output sequence released as output by the enforcer for φ after reading input-output sequence $\sigma \in \Sigma^*$. Upon reading a new event (x, y) , if what has been already computed as output by the enforcer $E_\varphi(\sigma)$ followed by (x, y) does not allow to satisfy the property φ , then the enforcer edits (x, y) using functions editl_{φ_1} and editO_φ . When the current event (x, y) has to be edited, note that there may be several possible solutions. For example, consider the property S_1 introduced in Example 1. Let $\sigma = (10, 1) \cdot (01, 0)$, and the output of the enforcer after processing σ will be $E_\varphi(\sigma) = (10, 1) \cdot (01, 0)$. Let the new event be $(11, 0)$, and $E_\varphi(\sigma) \cdot (11, 0) \not\models \varphi$, and the enforcer has to edit the new event $(11, 0)$. Note that $E_\varphi(\sigma)_I = 10 \cdot 01$, and $\text{editl}_{\varphi_1}(10 \cdot 01) = \{00, 01, 10\}$ and the enforcer can choose any element from $\text{editl}_{\varphi_1}(10 \cdot 01)$ as the transformed input.

Remark 3 (Enforcing bi-directional properties). By considering two uni-directional enforcers, where one enforcer checks and transforms inputs from the environment to the program and another enforcer checks and transforms outputs from the program to the environment, bi-directional properties cannot be enforced. For example, bi-directional properties such as the property S_1 introduced in Example 1 cannot be enforced using two uni-directional enforcers.

Remark 4 (When the input word provided to the enforcer satisfies φ). Constraint (\mathbf{Tr}') expresses that when any input-output word $\sigma \in \Sigma^*$ provided as input to the enforcer satisfies the property φ , then the enforcer will not edit any event and will output σ (i.e., $E_\varphi(\sigma) = \sigma$).

$$\forall \sigma \in \Sigma^* : E_\varphi(\sigma) \models \varphi \implies E_\varphi(\sigma) = \sigma. \quad (\mathbf{Tr}')$$

Lemma 2. $(\mathbf{Tr}) \implies (\mathbf{Tr}')$.

Lemma 2 shows that (\mathbf{Tr}') is a consequence of constraint (\mathbf{Tr}) . For any φ , for any $\sigma \in \Sigma^$, proof of this lemma is straightforward using induction on σ .*

σ	$E_\varphi(\sigma)$	\mathbf{Tr}	\mathbf{Tr}'
$(10, 1)$	$(10, 1)$	✓	✓
$(10, 1) \cdot (11, 1)$	$(10, 1) \cdot (\mathbf{10}, \mathbf{1})$	✓	✓
$(10, 1) \cdot (11, 1) \cdot (01, 0)$	$(10, 1) \cdot (\mathbf{10}, \mathbf{1}) \cdot (\mathbf{10}, \mathbf{0})$	✗	✓
$(10, 1) \cdot (11, 1) \cdot (01, 0)$	$(10, 1) \cdot (\mathbf{10}, \mathbf{1}) \cdot (01, 0)$	✓	✓

Table 1: Example: (\mathbf{Tr}) Vs. (\mathbf{Tr}')

Example 2 ((\mathbf{Tr}) is stronger than (\mathbf{Tr}')). Via this example, we illustrate that constraint (\mathbf{Tr}) is stronger than the alternative transparency constraint (\mathbf{Tr}') . Let us consider the property S_1 introduced in Example 1. In Table 1, first column denoted using σ shows input-output words, and the second column denoted using $E_\varphi(\sigma)$ shows the output of the enforcer for σ , and the next two columns indicate whether $E_\varphi(\sigma)$ satisfies constraints (\mathbf{Tr}) and (\mathbf{Tr}') respectively. We can see that there are situations where (\mathbf{Tr}') holds and (\mathbf{Tr}) does not hold. When the enforcer reads the third event $(01, 0)$, if it edits this event to $(10, 0)$, then constraint (\mathbf{Tr}') holds, and constraint (\mathbf{Tr}) does not hold since $E_\varphi((10, 1) \cdot (11, 1))$ followed by the new event read $(01, 0)$ satisfies the property S_1 , and it should not be edited by the enforcer according to constraint (\mathbf{Tr}) .

Definition 4 (Enforceability). Let $\varphi \subseteq \Sigma^*$ be a property. We say that φ is enforceable iff an enforcer E_φ for φ exists according to Definition 3.

Not all properties are enforceable, even if we restrict ourselves to prefix-closed safety properties, as the following example shows.




Fig. 3: A non-enforceable safety property.

Example 3 (Non-enforceable safety property). We illustrate that not all prefix-closed safety properties are enforceable according to Definition 3. Consider the automaton in Figure 3 defining the property φ that we want to enforce, with $I = \{A\}$, $O = \{B\}$ and $\Sigma = \Sigma_I \times \Sigma_O$. Let the input-output sequence provided as input to the enforcer be $\sigma = (1, 1) \cdot (1, 0)$. When the enforcer reads the first event $(1, 1)$, it can output $(1, 1)$ (since every event in Σ from q_0 leads to a non violating state q_1). Note that from q_1 , every event in Σ only leads to violating state q_v . Thus, when the second event $(1, 0)$ is read, every possible editing of this event will only lead to violation of the property. Upon reading the second event $(1, 0)$, releasing any event in Σ as output will violate soundness, and if no event is released as output, then the instantianey constraint will be violated.

Theorem 1 (Condition for enforceability). Consider a property φ defined as SA $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$. Property φ is enforceable iff the following condition holds:

$$\forall q \in Q, q \neq q_v \implies \exists (x, y) \in \Sigma : q \xrightarrow{(x, y)} q' \wedge q' \neq q_v \quad (\mathbf{EnfCo})$$

Proof of Theorem 1 is given in Appendix A, page 13. Note that given any property φ defined as SA $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$, it is straightforward to test whether \mathcal{A}_φ satisfies condition **(EnfCo)**.

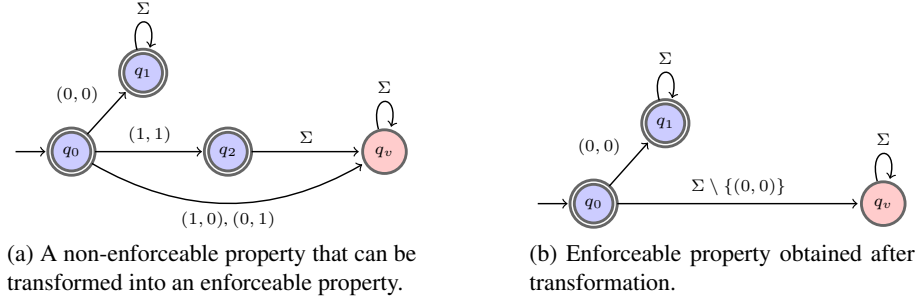


Fig. 4: A non-enforceable property transformed into an enforceable property.

Remark 5 (Transforming a non-enforceable property into an enforceable property). Some non-enforceable properties can be made enforceable by a transformation that excludes some behaviors from the property. We illustrate this with an example. Consider the property defined by the automaton in Figure 4a. This property is not enforceable for the following reason. Suppose that the first input-output event that the enforcer receives is $(1, 1)$. Since there is a transition from q_0 to q_2 upon $(1, 1)$, the enforcer will take this transition (according to transparency constraint). Then, whatever may be the second event that the enforcer receives, note that editI_{φ_1} and editO_{φ} will be empty, and there is no way to correct the event and avoid reaching q_v . However, we can transform this property into an enforceable property by excluding all the paths/behaviours that are problematic. In particular, we can remove state q_2 from the automaton of Figure 4a and redirect the transition labeled $(1, 1)$ from q_0 to q_v instead. This has the effect of removing the word $(1, 1)$ from the language accepted by this automaton. The resulting automaton (shown in Figure 4b) that we obtain satisfies the condition for enforceability **(EnfCo)** and therefore the resulting new property is enforceable. Note that transforming a non-enforceable property to an enforceable one is not always possible. For instance, the non-enforceable property of Figure 3 cannot be transformed to an enforceable property.

Transformation of non-enforceable properties. If a given safety property φ defined as SA $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$ does not satisfy the condition for enforceability **(EnfCo)**, then we can apply the following transformation process to check whether \mathcal{A}_φ can be transformed in to an enforceable property (by discarding some states in $Q \setminus \{q_0\}$ in the automaton \mathcal{A}_φ). We discuss the algorithm for transformation briefly.

- For every state $q \in Q \setminus \{q_v\}$ if $\forall (x, y) \in \Sigma, q \xrightarrow{(x, y)} q_v$, then merge q with q_v (q is removed from the set of states Q and all the incoming transitions to q go to q_v instead).
- The transformation continues until one of the following two conditions hold:
 - only two states q_0 and q_v remain in Q , i.e., $Q = \{q_0, q_v\}$ such that $\forall (x, y) \in \Sigma, q_0 \xrightarrow{(x, y)} q_v$. In this case, the algorithm returns that \mathcal{A}_φ cannot be transformed into an enforceable property.
 - $Q \setminus \{q_0, q_v\}$ is non-empty, and there is no state in $Q \setminus \{q_0, q_v\}$, that has all its outgoing transitions to q_v . In this case, the algorithm returns the resulting transformed automaton which is an enforceable property. Let $\text{sub}(\mathcal{A}_\varphi)$ be the transformed automaton. Note that $\mathcal{L}(\text{sub}(\mathcal{A}_\varphi)) \subseteq \mathcal{L}(\mathcal{A}_\varphi)$.

The algorithm for transformation of non-enforceable properties is discussed in detail in Appendix B.

4 Algorithm

In this section, we provide an algorithm for implementing the bi-directional synchronous enforcement problem defined in Section 3. Let the SA $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$ define the property φ that we want to enforce. SA $\mathcal{A}_{\varphi_I} = (Q, q_0, q_v, \Sigma_I, \rightarrow_I)$ is obtained from \mathcal{A}_φ by projecting on inputs (see section 2).

Algorithm 1 Enforcer

```

1:  $t \leftarrow 0$ 
2:  $q \leftarrow q_0$ 
3: while true do
4:    $x_t \leftarrow \text{read\_in\_chan}()$ 
5:   if  $\exists q' \in Q : q \xrightarrow{x_t}_I q' \wedge q' \neq q_v$  then
6:      $x'_t \leftarrow x_t$ 
7:   else
8:      $x'_t \leftarrow \text{nondet\_editl}_{\mathcal{A}_{\varphi_I}}(q)$ 
9:   end if
10:   $\text{ptick}(x'_t)$ 
11:   $y_t \leftarrow \text{read\_out\_chan}()$ 
12:  if  $\exists q' \in Q : q \xrightarrow{(x'_t, y_t)} q' \wedge q' \neq q_v$  then
13:     $y'_t \leftarrow y_t$ 
14:  else
15:     $y'_t \leftarrow \text{nondet\_editO}_{\mathcal{A}_\varphi}(q, x'_t)$ 
16:  end if
17:   $\text{release}((x'_t, y'_t))$ 
18:   $q \leftarrow q'$  where  $q \xrightarrow{(x'_t, y'_t)} q' \wedge q' \neq q_v$ 
19:   $t \leftarrow t + 1$ 
20: end while

```

We provide an online algorithm that requires automata \mathcal{A}_φ and \mathcal{A}_{φ_I} as input. Algorithm 1 is an infinite loop, and an iteration of the algorithm is triggered at every time step. We adapt the *reactive interface* that is used for linking the program to its adjoining environment by following the structure of the interface described in [2]. We extend the interface by including the enforcer as an intermediary between the synchronous program and its adjoining environment.

In the algorithm shown below, t keeps track of the time-step (*tick*), initialized with 0. q keeps track of the current state of both the automata \mathcal{A}_φ and \mathcal{A}_{φ_I} . Recall that the automaton \mathcal{A}_{φ_I} that we obtain from the automaton \mathcal{A}_φ by projecting on inputs (see Section 2) have identical structure, and the only difference is that the outputs are ignored on the transitions in the automaton \mathcal{A}_{φ_I} . Note that at the beginning of each iteration of the algorithm, the current states of both the automata \mathcal{A}_φ and \mathcal{A}_{φ_I} are the same (where both are initialized with q_0). At t , if $\text{EOut} \in \Sigma^*$ is the input-output sequence obtained by concatenating all the events released as output by the enforcer until time t , then q corresponds to the state that we reach in the automaton \mathcal{A}_φ upon reading EOut . Similarly, if $\text{EOut}_I \in \Sigma_I^*$ is the sequence obtained by projecting on x'_t s from EOut , q also corresponds to the state that we reach in the automaton \mathcal{A}_{φ_I} upon reading EOut_I .

Functions read_in_chan (resp. read_out_chan) are functions corresponding to reading input (resp. output) channels, and function ptick corresponds to invoking the synchronous program. Function release takes an input-output event, and releases it as output of the enforcer.

Each iteration of the algorithm proceeds as follows: first all the input channels are read using function read_in_chan and the input event is assigned to x_t . Then the algorithm tests whether there exists a transition in \rightarrow_I from the current state q upon x_t to an accepting state in \mathcal{A}_{φ_I} . In case if this test succeeds, then it is not necessary to edit the input event x_t , and the transformed input x'_t is assigned x_t . Otherwise, x'_t is assigned with the output of $\text{nondet_editl}_{\mathcal{A}_{\varphi_I}}(q)$. Let us recall that $\text{nondet_editl}_{\mathcal{A}_{\varphi_I}}(q)$ returns an input event that leads to an accepting state in \mathcal{A}_{φ_I} from q .

After transforming the input x_t according to \mathcal{A}_{φ_I} , the program is invoked with the transformed input x'_t using function `ptick`. Afterwards, all the output channels are read using function `read_out_chan` and the output event is assigned to y_t . Then the algorithm tests whether there exists a transition in \rightarrow from the current state q upon (x'_t, y_t) to an accepting state in \mathcal{A}_{φ} . In case if this test succeeds, then it is not necessary to edit the output event y_t , and the transformed output y'_t is assigned y_t . Otherwise, y'_t is assigned with the output of `nondet-editO $_{\mathcal{A}_{\varphi}}$` (q, x'_t). Note that `nondet-editO $_{\mathcal{A}_{\varphi}}$` (q, x'_t) returns an output event y'_t such that (x'_t, y'_t) leads to an accepting state in \mathcal{A}_{φ} from q .

Before proceeding with the next iteration, current state q is updated to q' which is the state reached upon (x'_t, y'_t) from state q in the automaton \mathcal{A}_{φ} , and the time-step t is incremented. Note that if there exists a transition $q \xrightarrow{(x'_t, y'_t)} q'$ in the SA \mathcal{A}_{φ} , then there also exists a transition $q \xrightarrow{x'_t}_I q'$ in the SA \mathcal{A}_{φ_I} . The current states of both the SA are always synchronized and the same at the beginning of each iteration of the algorithm.

Definition 5 (E_{φ}^*). Consider an enforceable safety property φ . We define the function $E_{\varphi}^* : \Sigma^* \rightarrow \Sigma^*$, where $\Sigma = \Sigma_I \times \Sigma_O$, as follows. Let $\sigma = (x_1, y_1) \cdots (x_k, y_k) \in \Sigma^*$ be a word received by Algorithm 1. Then we let $E_{\varphi}^*(\sigma) = (x'_1, y'_1) \cdots (x'_k, y'_k)$, where (x'_t, y'_t) is the pair of events output by Algorithm 1 in Step 17, for $t = 1, \dots, k$.

Theorem 2 (Correctness of the enforcement algorithm). Given any safety property φ defined as SA \mathcal{A}_{φ} that satisfies condition **(EnfCo)**, the function E_{φ}^* defined above is an enforcer for φ , that is, it satisfies **(Snd)**, **(Tr)**, **(Mono)**, **(Inst)**, and **(Cau)** constraints of Definition 3.

Proof of Theorem 2 is given in Appendix A, page 13.

Remark 6 (Determinism of the enforcer). Since we consider synchronous programs, the enforcer should be deterministic. Regarding determinism, note that though \mathcal{A}_{φ} is deterministic, the enforcer E_{φ}^* may be non-deterministic, because when the received input x (resp. output y) does not lead to an accepting state from the current state q in \mathcal{A}_{φ_I} (resp. \mathcal{A}_{φ}), it is edited in step 8 (resp. step 15) of the algorithm. Note that `editl $_{\mathcal{A}_{\varphi_I}}$` (q) (resp. `editO $_{\mathcal{A}_{\varphi}}$` (q, x) where $x \in \text{editl}_{\mathcal{A}_{\varphi_I}}(q)$), may contain more than one element as illustrated via an example in Remark 2, and `nondet-editl $_{\mathcal{A}_{\varphi_I}}$` (resp. `nondet-editO $_{\mathcal{A}_{\varphi}}$`) will choose one element from the set `editl $_{\mathcal{A}_{\varphi_I}}$` (q) (resp. `editO $_{\mathcal{A}_{\varphi}}$` (q, x)). However, it is straightforward to make the behavior deterministic by computing `editl $_{\mathcal{A}_{\varphi_I}}$` (q) off-line for all $q \in Q \setminus \{q_v\}$, and selecting one element randomly from `editl $_{\mathcal{A}_{\varphi_I}}$` (q) and remembering the selection for each q by storing in a table with size $|Q|$. Thus, whenever in some state q and when the input read x does not lead to an accepting state from q (i.e, the condition tested in line 5 evaluates to false), in step 8 we check the element corresponding to the state q from the table and assign it to x' . Similarly, for event $q \in Q \setminus \{q_v\}$, and for all $x \in \text{editl}_{\mathcal{A}_{\varphi_I}}(q)$, we can compute `editO $_{\mathcal{A}_{\varphi}}$` (q, x) off-line, select one element randomly and store the selection in a table with size $|Q \times \Sigma_I|$. Thus, whenever in some state q , when (x', y) (where x' is the transformed input and y is the output read) does not lead to an accepting state, in step 15 we check the element corresponding to (q, x') from the table and assign it to y' .

5 Application to SCCharts

SCCharts is a Statechart dialect that has been designed for safety-critical systems and offers deterministic concurrency [21]. We implemented the algorithm presented in Section 4 in an SCCharts compilation framework⁵ according to the *single-pass language-driven incremental compilation* approach [14]. Here, a safety automaton is automatically transformed into a synchronous enforcer using model-to-model transformations. The generated enforcer has three concurrent regions, one for reading and editing the inputs, one for invoking the tick function `ptick` with the edited inputs and a final one for processing and emitting the outputs. The three components exactly match the steps of the algorithm presented in Section 4.

Figure 5 depicts the example safety automaton ABO SA in SCCharts and the automatically generated Enforcer ABO Enf. In this example, A and B serve as input vector, whereas O is the only output. The automaton only has two states, the initial state q_0 and the violation state q_v . The safety property says that A and B and also B and O may not be present at the same time.

⁵ <https://rtsys.informatik.uni-kiel.de/kieler>

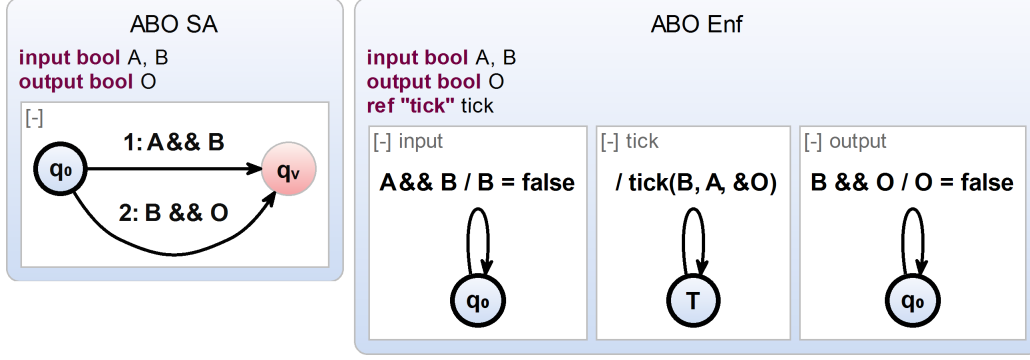


Fig. 5: Example safety automaton ABO in SCCharts (left) and its automatically generated enforcer (right).

Examples ⁶	Tick (LoC)	# Properties	Enf. (LoC)	Time (μs)	Time w/ Enf. (μs)	Incr. (%)
Null	0	0	0	0.654	0.752	14.98
ABRO	23	1	21	1.208	1.565	29.55
ABO	28	1	21	0.998	1.368	37.10
Reactor	32	2	32	1.587	2.137	34.61
Faulty Heart Model	43	2	40	1.346	1.869	38.85
Simple Heart Model	76	2	40	2.175	2.825	29.86
Traffic Light	171	3	41	4.039	4.707	16.53
Pacemaker	271	2	35	7.302	8.318	13.91
FHM + Pacemaker	314	2	35	9.195	10.306	12.08

Table 2: Evaluation results.

Remark 7. In the ABO example, Figure 5, two regions (tick and output) concurrently write to a shared variable O . Usually, this would be considered a write-write race, leading either to non-determinism, e.g. in Java threads, or rejection at compile time due to non-causality [3], as in synchronous languages. However, we can take advantage of the fixed execution sequence of the three regions during every tick, following Algorithm 1. First a transition in the input region is executed, followed by the tick region and finally the output region. This approach, thus, follows the PRET-C [1] semantics, which is causal by construction.

In order to evaluate this implementation we used a series of models with increasing sizes as can be seen in Table 2. To generate the mean values we simulated every model 5 times with each run consisting of 1000 ticks. As inputs for each model, a random environment was created. The whole setup was executed for two cases. Firstly, the plain model was simulated within its environment. Secondly, the same environment was used to simulate the model again with an enforcer in between. The number of enforced properties (entry “# Properties” in Table 2) range from 0 to 3 including properties that enforce inputs and also outputs (bi-directional properties). All experiments were conducted on an embedded system equipped with an 1 GHz ARM Cortex-A7 Dual-Core. Depending on the model size and the number of properties enforced, we see an increase of mean execution time between 12%-38% when simulating with an enforcer. Due to the netlist-based code generation of KIELER, there is a constant overhead because of the tick function call. Therefore, the overhead decreases percentage-wise with increasing model size. The Null model test measures the overhead of this black-box call with an enforcer with 0 safety properties. We observe a constant overhead of $0.1 \mu s$ here.

² ABRO from [4], ABO from [21], Reactor from [20], Simple Heart Model and Pacemaker are remodeled SCCharts variants from [10], Faulty Heart Model is a variant of the Simple Heart Model with deliberately flawed pulse signals, Traffic Light from [13] (remodeled from Ptolemy Traffic Light)

As a concrete case study, we selected a pacemaker based on [10], which has been implemented in SCCharts. As second experiment we ran the **Faulty Heart Model** together with the **Pacemaker**. The results of the close-loop simulation can be seen in the last row of Table 2. Here, the **Faulty Heart Model** serves as environment for the **Pacemaker** and generates flawed pulse signals for the heart. We added an enforcer to the pacemaker to make sure that atrial and ventricular signals cannot occur simultaneously, which results in editing the input vector, and also that the pacemaker does not emit pace signals for both in return, which results in editing the output vector. We observe a mean overhead of 12% when using the enforcer.

6 Related Work

Synthesizing enforcers from properties is an active area of research. According to how an enforcer is allowed to correct the input sequence, several RE models have been proposed. Security automata proposed by Schneider [19] focus on enforcement of safety properties, where the enforcer blocks the execution when it recognizes a sequence of actions that does not satisfy the desired property. Edit automata [12] allows the enforcer to correct the input sequence by suppressing and (or) inserting events, and the RE mechanisms proposed in [8,15] allows buffering events and releasing them upon observing a sequence that satisfies the desired property. Recently, compositionality of enforcers has been studied in [16]. Given a set of properties over the same alphabet, the problem studied in [16] addresses whether it is possible to synthesize multiple enforcers, one for each property, and whether composing enforcers (in series or in parallel) can enforce all the properties. Moreover, the enforcement framework in [16] allows to buffer (delay) events. These approaches focus on uni-directional RE.

Mandatory Result Automata (MRAs) [6] extended edit-automata [12], by considering bi-directional runtime enforcement. Compared to the other RE frameworks such as [19,8,12,15], in MRA the focus is on handling communication between two parties. However none of the above approaches are suitable for reactive systems since halting the program and delaying actions is not suitable. This is because for reactive systems the enforcer has to react instantaneously.

Our work is closely related to [5], which introduces a framework to synthesize enforcers for reactive systems, called as *shields*, from a set of safety properties. In our work, we restrict to prefix-closed safety properties. The approach in [5] seems to consider more that prefix-closed properties (where properties are expressed as automata), but not all regular properties. Also, the approach in [5] has the notion of k -stabilization where the shield allows to deviate from the property for k consecutive steps whenever a property violation is unavoidable. If a second violation occurs within k steps, then the shield enters into a *fail-safe* mode, where it ensures only correctness. So, if two or more errors occur within k -steps, then the shield may generate outputs arbitrarily to satisfy the property being monitored by ignoring outputs from the system being monitored. In our approach, if the input given to the enforcer satisfies the property, then the enforcer does not modify any event. In case if a violation is noticed upon some event, the enforcer corrects it (to avoid violation), and continues to minimize deviation also for the future input events depending on the state of the enforcer and the received input event. Moreover, in [5], the shield is uni-directional, where it observes inputs from the environment and outputs from the system (program), and transforms erroneous outputs. In our work, we consider bi-directional enforcement, as explained and illustrated in Fig. 1.

Note that when we consider safety-critical embedded systems such as medical devices and automotive systems, it is also utmost important to monitor and transform “illegal” inputs, before they are fed to the program. For instance, suppose that there are multiple sensors, and their values are inputs from the environment to the enforcer. Some sensors may fail or may be attacked by some intruder. Unlike [5], which ignores inconsistent inputs, our work is able to deal with both the inputs (from the environment) and the outputs (from the synchronous program) simultaneously during each reaction.

7 Conclusions

Synchronous observers are used to express safety properties for synchronous programs, which may be verified either statically or during runtime. This paper extends observers by proposing the concept of runtime enforcers for synchronous programs. The property to be enforced is modeled as a safety automaton, which is syntactically like

an observer (expressed as an automaton with a single violation state) referring to both inputs and outputs of the synchronous program. We formalise, for the first time, the runtime enforcement synthesis problem for synchronous reactive systems. We define enforceability conditions, provide an algorithm, and prove its correctness. The synthesised enforcer interacts with a black-box synchronous program and its adjoining environment to ensure that the property in question holds during program execution. We have implemented the proposed enforcer synthesis algorithm for the SCCharts synchronous language. We highlight the applicability of the proposed approach by enforcing policies over a synchronous pacemaker model. In the near future, we will consider several extensions, including enforcement with valued inputs and outputs (valued signals), non-safety properties, and distributed enforcement.

References

1. S. Andalam, P. S. Roop, A. Girault, and C. Traulsen. A predictable framework for safety-critical embedded systems. *IEEE Transactions on Computers*, 63(7):1600–1612, 2014.
2. C. Andre, F. Boulanger, and A. Girault. Software implementation of synchronous programs. In *Application of Concurrency to System Design, 2001. Proceedings. 2001 International Conference on*, pages 133–142. IEEE, 2001.
3. A. Benveniste, P. Caspi, S. Edwards, N. Halbwachs, P. Le Guernic, and R. de Simone. The synchronous languages 12 years later. *Proceedings of the IEEE*, 91(1):64–83, Jan 2003.
4. G. Berry. *The Esterel v5 Language Primer, Version v5_91*. Centre de Mathématiques Appliquées Ecole des Mines and INRIA, 06565 Sophia-Antipolis, 2000.
5. R. Bloem, B. Könighofer, R. Könighofer, and C. Wang. Shield synthesis: Runtime enforcement for reactive systems. In *TACAS*, volume 9035 of *LNCs*. Springer, 2015.
6. E. Dolzhenko, J. Ligatti, and S. Reddy. Modeling runtime enforcement with mandatory results automata. *Int. J. Inf. Sec.*, 14(1):47–60, 2015.
7. Y. Falcone, K. Havelund, and G. Reger. A tutorial on runtime verification. In *Engineering Dependable Software Systems*, volume 34, pages 141–175. IOS Press, 2013.
8. Y. Falcone, L. Mounier, J.-C. Fernandez, and J.-L. Richier. Runtime enforcement monitors: composition, synthesis, and enforcement abilities. *FMSD*, 38(3):223–262, 2011.
9. N. Halbwachs, F. Lagnier, and P. Raymond. Synchronous observers and the verification of reactive systems. In *Algebraic Methodology and Software Technology (AMAST93)*, pages 83–96. Springer, 1994.
10. Z. Jiang, M. Pajic, S. Moarref, R. Alur, and R. Mangharam. Modeling and verification of a dual chamber implantable pacemaker. In *TACAS*, pages 188–203. Springer, 2012.
11. M. Leucker and C. Schallhart. A brief account of runtime verification. *Journal of Logic and Algebraic Programming*, 78(5):293–303, 2009.
12. J. Ligatti, L. Bauer, and D. Walker. Run-time enforcement of nonsafety policies. *ACM Trans. Inf. Syst. Secur.*, 12(3):19:1–19:41, Jan. 2009.
13. C. Motika, H. Fuhrmann, R. von Hanxleden, and E. A. Lee. Executing domain-specific models in Eclipse. Technical Report 1214, Christian-Albrechts-Universität zu Kiel, Department of Computer Science, October 2012. ISSN 2192-6247.
14. C. Motika, S. Smyth, and R. von Hanxleden. Compiling SCCharts—A case-study on interactive model-based compilation. In *ISoLA*, volume 8802 of *LNCs*, pages 443–462, Corfu, Greece, 2014.
15. S. Pinisetty, Y. Falcone, T. Jéron, H. Marchand, A. Rollet, and O. Nguena Timo. Runtime enforcement of timed properties revisited. *FMSD*, 45(3):381–422, 2014.
16. S. Pinisetty and S. Tripakis. Compositional runtime enforcement. In *NASA Formal Methods Symposium, NFM 2016, Minneapolis, MN, USA*, pages 82–99. Springer, 2016.
17. P. Raymond, X. Nicollin, N. Halbwachs, and D. Weber. Automatic testing of reactive systems. In *Real-Time Systems Symposium*, pages 200–209. IEEE, 1998.
18. J. Rushby. The versatile synchronous observer. In *Specification, Algebra, and Software*, pages 110–128. Springer, 2014.
19. F. B. Schneider. Enforceable security policies. *ACM Trans. Inf. Syst. Secur.*, 3(1):30–50, 2000.
20. C. Traulsen, T. Amende, and R. von Hanxleden. Compiling SyncCharts to Synchronous C. In *Proceedings of the Design, Automation and Test in Europe Conference (DATE’11)*, pages 563–566, Grenoble, France, March 2011. IEEE.
21. R. von Hanxleden, B. Duderstadt, C. Motika, S. Smyth, M. Mandler, J. Aguado, S. Mercer, and O. O’Brien. SCCharts: Sequentially constructive statecharts for safety-critical applications. In *PLDI*, pages 372–383, NY, USA, 2014. ACM.

A Appendix: Proofs

Proof (of Theorem 1). Let us recall Theorem 1. Consider a property φ defined as SA $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)^7$. Property φ is enforceable iff the condition **(EnfCo)** holds which is the following condition: $\forall q \in Q, q \neq q_v \implies \exists (x, y) \in \Sigma : q \xrightarrow{(x, y)} q' \wedge q' \neq q_v$.

We prove that:

- *Sufficient:* If condition **(EnfCo)** holds then E_φ according to Definition 3 exists.

Due to condition **(EnfCo)**, whatever may be the current state $q \in Q \setminus \{q_v\}$ of the enforcer, there is at least one possibility to correct the event that it receives when in state q (in case if the received event leads to q_v from q). That is, due to condition **(EnfCo)**, $\forall q \in Q \setminus \{q_v\}$, we know for sure that $\text{editl}_{\mathcal{A}_{\varphi_1}}(q)$ will be non-empty, and $\forall q \in Q \setminus \{q_v\}, \forall x \in \text{editl}_{\mathcal{A}_{\varphi_1}}(q) : \text{editO}_{\mathcal{A}_\varphi}(q, x)$ will be also non-empty.

For any property φ defined as SA \mathcal{A}_φ , the enforcement function E_φ^* (Definition 5) is an enforcer for φ which satisfies all the constraints according to Definition 3. Theorem 2 shows that for any property φ (defined as SA \mathcal{A}_φ) that satisfies the condition for enforceability **(EnfCo)**, the enforcement function E_φ^* (Definition 5) is an enforcer for φ , that is, it satisfies **(Snd)**, **(Tr)**, **(Mono)**, **(Inst)**, and **(Cau)** constraints of Definition 3.

- *Necessary:* If E_φ according to Definition 3 exists, then condition **(EnfCo)** holds.

Suppose that an enforcer E_φ for φ according to Definition 3 exists and assume that condition **(EnfCo)** does not hold for \mathcal{A}_φ .

Since condition **(EnfCo)** does not hold, $\exists q \in Q \setminus \{q_v\} : \forall (x, y) \in \Sigma : q \xrightarrow{(x, y)} q_v$, i.e., there exists a location $q \in Q \setminus \{q_v\}$ such that all the outgoing transitions from q go to q_v .

Since all the locations in Q are reachable from q_0 , $\exists \sigma \in \Sigma^* : q_0 \xrightarrow{\sigma} q$, i.e., there certainly exists a word $\sigma \in \Sigma^*$ that leads to the problematic accepting location q (which has all its outgoing transitions to q_v) from the initial location q_0 .

If σ is the input word to the enforcer, then due to constraint **(Tr)**, it cannot edit any event in σ , and the enforcer produces σ as output and reaches location q . When in location q , upon receiving any event $(x, y) \in \Sigma$, the enforcer has no possibility to correct it, since every event in Σ leads to q_v from q (i.e., since $\forall (x, y) \in \Sigma : q \xrightarrow{(x, y)} q_v$).

Thus, $\forall (x, y) \in \Sigma$, when the input word given to the enforcer is $\sigma \cdot (x, y)$, the enforcer cannot produce any event as output since $\text{editl}_{\varphi_1}()$ and $\text{editO}_\varphi()$ from location q will be empty, violating constraints **(Inst)** and **(Cau)**. Thus, our assumption is false and condition **(EnfCo)** holds for \mathcal{A}_φ .

Proof (of Theorem 2). Let us recall the condition for enforceability: A property φ defined as SA $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$ is enforceable iff

$$\forall q \in Q, q \neq q_v \implies \exists (x, y) \in \Sigma : q \xrightarrow{(x, y)} q' \wedge q' \neq q_v.$$

Let us also recall the definition of function $E_\varphi^* : \Sigma^* \rightarrow \Sigma^*$ (Definition 5). Let $\sigma = (x_1, y_1) \cdots (x_k, y_k) \in \Sigma^*$ be a word received by Algorithm 1. Then we let $E_\varphi^*(\sigma) = (x'_1, y'_1) \cdots (x'_k, y'_k)$, where (x'_t, y'_t) is the pair of events output by Algorithm 1 in Step 17, for $t = 1, \dots, k$.

Note that the input automaton $\mathcal{A}_{\varphi_I} = (Q, q_0, q_v, \Sigma_I, \rightarrow_I)$ is obtained from \mathcal{A}_φ by projecting on inputs (See Definition 2, Section 2).

We shall prove that given any safety property φ defined as SA \mathcal{A}_φ that satisfies condition **(EnfCo)**, the function E_φ^* is an enforcer for φ , that is, it satisfies **(Snd)**, **(Tr)**, **(Mono)**, **(Inst)**, and **(Cau)** constraints of Definition 3.

Let us prove this theorem using induction on the length of the input sequence $\sigma \in \Sigma^*$ (which also corresponds to the number of ticks/iterations of Algorithm 1).

Induction basis. Theorem 2 holds trivially for $\sigma = \epsilon$ since the algorithm will not release any input-output event as output and thus $E_\varphi^*(\epsilon) = \epsilon$.

Induction step. Assume that for every $\sigma = (x_1, y_1) \cdots (x_k, y_k) \in \Sigma^*$ of some length $k \in \mathbb{N}$, let $E_\varphi^*(\sigma) = (x'_1, y'_1) \cdots (x'_k, y'_k) \in \Sigma^*$, for $t = 1, \dots, k$, and Theorem 2 holds for σ , i.e., $E_\varphi^*(\sigma)$ satisfies the **(Snd)**, **(Tr)**, **(Mono)**, **(Inst)**, and **(Cau)** constraints. Let $q \in Q \setminus \{q_v\}$ be the current state of both the automata \mathcal{A}_φ and \mathcal{A}_{φ_I} after processing

⁷ Note that we consider that \mathcal{A}_φ is deterministic and complete, and Q does not contain any (redundant) locations that are not reachable from q_0 in 1 or more steps.

input σ of length k , i.e., q corresponds to the state that we reach upon $E_\varphi^*(\sigma)$ in \mathcal{A}_φ , and the state that we reach in the automaton \mathcal{A}_{φ_I} upon $E_\varphi^*(\sigma)_I$. Note that the current state q in Algorithm 1 can never be q_v (q is initialized to q_0 and it is updated in step 18 to a state $q' \in Q \setminus \{q_v\}$).

We now prove that for any event $(x_{k+1}, y_{k+1}) \in \Sigma$, Theorem 2 holds for $\sigma \cdot (x_{k+1}, y_{k+1})$, where $x_{k+1} \in \Sigma_I$ is the input event read by Algorithm 1, and $y_{k+1} \in \Sigma_O$ is the output event read by Algorithm 1 in $k+1^{th}$ iteration (i.e., when $t = k+1$). We have the following two possible cases based on whether there is a transition in the automaton \mathcal{A}_φ from the current state q upon (x_{k+1}, y_{k+1}) to an accepting state.

- $\exists q' \in Q : q \xrightarrow{(x_{k+1}, y_{k+1})} q' \wedge q' \neq q_v$.

In Algorithm 1, the condition tested in step 5 will evaluate to true since from Lemma 1, in \mathcal{A}_{φ_I} we will have

$$\exists q' \in Q : q \xrightarrow{x_{k+1}}_I q' \wedge q' \neq q_v, \text{ and thus } x'_{k+1} = x_{k+1}.$$

Also, the condition tested in step 12 will evaluate to true in this case since $\exists q' \in Q : q \xrightarrow{(x_{k+1}, y_{k+1})} q' \wedge q' \neq q_v$, and thus $y'_{k+1} = y_{k+1}$. At the end of the $k+1^{th}$ iteration, the input-output event released as output by the algorithm in step 17 is (x_{k+1}, y_{k+1}) . The output of the algorithm after completing the $k+1^{th}$ iteration is $E_\varphi^*(\sigma \cdot (x_{k+1}, y_{k+1})) = E_\varphi^*(\sigma) \cdot (x_{k+1}, y_{k+1})$.

Regarding constraint **(Snd)**, in this case, what has been already released as output by the algorithm earlier before reading event (x_{k+1}, y_{k+1}) (i.e., $E_\varphi^*(\sigma)$) followed by the new input-output event released as output (x_{k+1}, y_{k+1}) satisfies the property φ , and thus constraint **(Snd)** holds.

Regarding constraint **(Mono)**, it holds since $\sigma \preceq \sigma \cdot (x_{k+1}, y_{k+1})$ and also $E_\varphi^*(\sigma) \preceq E_\varphi^*(\sigma) \cdot (x_{k+1}, y_{k+1})$.

Regarding constraint **(Inst)** from the induction hypothesis, we have for σ of some length k , $|\sigma| = |E_\varphi(\sigma)|$. We also have $E_\varphi^*(\sigma \cdot (x_{k+1}, y_{k+1})) = E_\varphi^*(\sigma) \cdot (x_{k+1}, y_{k+1})$. Thus, $|\sigma \cdot (x_{k+1}, y_{k+1})| = |E_\varphi^*(\sigma \cdot (x_{k+1}, y_{k+1}))| = k+1$, and constraint **(Inst)** holds.

Constraint **(Tr)** holds in this case since the output of the enforcer before reading (x_{k+1}, y_{k+1}) i.e., $E_\varphi^*(\sigma)$ followed by the new input-output event read (x_{k+1}, y_{k+1}) satisfies the property φ and we already saw that the output event released by the algorithm after reading (x_{k+1}, y_{k+1}) is $E_\varphi^*(\sigma) \cdot (x_{k+1}, y_{k+1})$.

Regarding constraint **(Cau)**, in this case from the induction hypothesis, from the definitions of $\text{editl}_{\mathcal{A}_{\varphi_I}}$ and $\text{editO}_{\mathcal{A}_\varphi}$ we have $x_{k+1} \in \text{editl}_{\mathcal{A}_{\varphi_I}}(q)$, and also $y_{k+1} \in \text{editO}_{\mathcal{A}_\varphi}(q, x_{k+1})$.

Theorem 2 thus holds for $\sigma \cdot (x_{k+1}, y_{k+1})$ in this case.

- $\nexists q' \in Q : q \xrightarrow{(x_{k+1}, y_{k+1})} q' \wedge q' \neq q_v$.

In this case, we have two sub-cases, based on whether $\exists q' \in Q : q \xrightarrow{x_{k+1}}_I q' \wedge q' \neq q_v$ in \mathcal{A}_{φ_I} .

- $\exists q' \in Q : q \xrightarrow{x_{k+1}}_I q' \wedge q' \neq q_v$.

In Algorithm 1, the condition tested in step 5 will evaluate to true and thus $x'_{k+1} = x_{k+1}$.

In this case, the condition tested in step 12 will evaluate to false since $\nexists q' \in Q : q \xrightarrow{(x_{k+1}, y_{k+1})} q' \wedge q' \neq q_v$. y'_{k+1} will thus be an element belonging to the set $\text{editO}_{\mathcal{A}_\varphi}(q, x_{k+1})$ if $\text{editO}_{\mathcal{A}_\varphi}(q, x_{k+1})$ is non-empty. It is important to notice that $\text{editO}_{\mathcal{A}_\varphi}(q, x_{k+1})$ will be non-empty in this case since we know for sure that

$\exists y'_{k+1} \in \Sigma_O, q' \in Q : q \xrightarrow{(x_{k+1}, y'_{k+1})} q' \wedge q' \neq q_v$ (from the condition for enforceability **(EnfCo)**, hypothesis ($q \neq q_v$), definition of $\text{editO}_{\mathcal{A}_\varphi}$, and Lemma 1). Thus y'_{k+1} is an element belonging to $\text{editO}_{\mathcal{A}_\varphi}(q, x_{k+1})$. The output of the algorithm after completing the $k+1^{th}$ iteration is $E_\varphi^*(\sigma \cdot (x_{k+1}, y_{k+1})) = E_\varphi^*(\sigma) \cdot (x_{k+1}, y'_{k+1})$. Regarding constraint **(Snd)**, from the definition of $\text{editO}_{\mathcal{A}_\varphi}$, we know that $E_\varphi^*(\sigma)$ followed by the new input-output event released as output (x_{k+1}, y'_{k+1}) satisfies property φ , and thus constraint **(Snd)** holds.

The reasoning for constraints **(Mono)** and **(Inst)** are similar to the previous cases since we saw that Algorithm 1 releases a new event (x_{k+1}, y'_{k+1}) as output after reading event (x_{k+1}, y_{k+1}) after completing $k+1^{th}$ iteration.

Constraint **(Tr)** holds trivially in this case since $E_\varphi^*(\sigma) \cdot (x_{k+1}, y_{k+1}) \not\models \varphi$.

Regarding constraint **(Cau)**, in this case from the induction hypothesis, from the definitions of $\text{editl}_{\mathcal{A}_{\varphi_I}}$ we have $x_{k+1} \in \text{editl}_{\mathcal{A}_{\varphi_I}}(q)$, and we already discussed that $\text{editO}_{\mathcal{A}_\varphi}(q, x_{k+1})$ will be non-empty and thus constraint **(Cau)** holds in this case.

- $\nexists q' \in Q : q \xrightarrow{x_{k+1}}_I q' \wedge q' \neq q_v$.

In Algorithm 1, the condition tested in step 5 will evaluate to false in this case. It is important to notice that $\text{editl}_{\mathcal{A}_{\varphi_1}}(q)$ will be non-empty since from the condition for enforceability and Lemma 1, we know for sure that $\exists x' \in \Sigma_I, q' \in Q : q \xrightarrow{x'_{k+1}} q' \wedge q' \neq q_v$ in the automaton \mathcal{A}_{φ_I} . Thus, x'_{k+1} will be an element belonging to $\text{editl}_{\mathcal{A}_{\varphi_1}}(q)$.

We have two sub-cases based on whether $\exists q' \in Q : q \xrightarrow{(x'_{k+1}, y_{k+1})} q' \wedge q' \neq q_v$ or not.

* $\exists q' \in Q : q \xrightarrow{(x'_{k+1}, y_{k+1})} q' \wedge q' \neq q_v$.

In Algorithm 1, the condition tested in step 12 will evaluate to true in this case. Thus, $y'_{k+1} = y_{k+1}$ in this case and the event released as output by the algorithm at the end of $k + 1^{th}$ iteration is (x'_{k+1}, y_{k+1}) . We have $E_{\varphi}^*(\sigma \cdot (x_{k+1}, y_{k+1})) = E_{\varphi}^*(\sigma) \cdot (x'_{k+1}, y_{k+1})$.

Regarding constraint (**Snd**), from the condition of this case (i.e., $\exists q' \in Q : q \xrightarrow{(x'_{k+1}, y_{k+1})} q' \wedge q' \neq q_v$), we know that $E_{\varphi}^*(\sigma)$ followed by the new input-output event released as output (x'_{k+1}, y_{k+1}) satisfies the property φ , and thus constraint (**Snd**) holds.

The reasoning for constraints (**Mono**) and (**Inst**) are similar to the previous cases since we saw that the algorithm releases a new event (x'_{k+1}, y_{k+1}) as output after reading event (x_{k+1}, y_{k+1}) at the end of $k + 1^{th}$ iteration.

Constraint (**Tr**) holds trivially in this case since $E_{\varphi}^*(\sigma) \cdot (x_{k+1}, y_{k+1}) \not\models \varphi$.

Regarding constraint (**Cau**), we already discussed that $\text{editl}_{\mathcal{A}_{\varphi_1}}(q)$ is non-empty and $x'_{k+1} \in \text{editl}_{\mathcal{A}_{\varphi_1}}(q)$, and $y_{k+1} \in \text{editO}_{\mathcal{A}_{\varphi}}(q, x'_{k+1})$ from the condition of this case and definitions of $\text{editl}_{\mathcal{A}_{\varphi_1}}$ and $\text{editO}_{\mathcal{A}_{\varphi}}$.

* $\nexists q' \in Q : q \xrightarrow{(x'_{k+1}, y_{k+1})} q' \wedge q' \neq q_v$.

In the algorithm, the condition tested in step 12 will evaluate to false in this case.

y'_{k+1} will thus be an element belonging to the set $\text{editO}_{\mathcal{A}_{\varphi}}(q, x'_{k+1})$ if $\text{editO}_{\mathcal{A}_{\varphi}}(q, x'_{k+1})$ is non-empty. Note that $\text{editO}_{\mathcal{A}_{\varphi}}(q, x'_{k+1})$ will be non-empty in this case since we know for sure that $\exists y'_{k+1} \in \Sigma_O, q' \in$

$Q : q \xrightarrow{(x'_{k+1}, y'_{k+1})} q' \wedge q' \neq q_v$ (from the enforceability condition, definitions, and Lemma 1). Thus y'_{k+1} is an element belonging to $\text{editO}_{\mathcal{A}_{\varphi}}(q, x'_{k+1})$. The output of the algorithm after completing the $k + 1^{th}$ iteration is $E_{\varphi}^*(\sigma \cdot (x_{k+1}, y_{k+1})) = E_{\varphi}^*(\sigma) \cdot (x'_{k+1}, y'_{k+1})$ where x'_{k+1} is an element belonging to $\text{editl}_{\mathcal{A}_{\varphi_1}}(q_I)$ and y'_{k+1} is an element belonging to $\text{editO}_{\mathcal{A}_{\varphi}}(q, x'_{k+1})$.

Regarding constraint (**Snd**), from the definitions of $\text{editl}_{\mathcal{A}_{\varphi_1}}$ and $\text{editO}_{\mathcal{A}_{\varphi}}$, we know that $E_{\varphi}(\sigma) \cdot (x'_{k+1}, y'_{k+1})$ satisfies the property φ and thus constraint (**Snd**) holds.

The reasoning for constraints (**Mono**) and (**Inst**) are similar to the previous cases since we saw that the algorithm releases a new event (x'_{k+1}, y'_{k+1}) as output after reading event (x_{k+1}, y_{k+1}) .

Constraint (**Tr**) holds trivially in this case since $E_{\varphi}(\sigma) \cdot (x_{k+1}, y_{k+1}) \not\models \varphi$.

Regarding constraint (**Cau**), we already discussed that $\text{editl}_{\mathcal{A}_{\varphi_1}}(q)$ is non-empty and $x'_{k+1} \in \text{editl}_{\mathcal{A}_{\varphi_1}}(q)$, and $\text{editO}_{\mathcal{A}_{\varphi}}(q, x'_{k+1})$ is also non-empty and $y'_{k+1} \in \text{editO}_{\mathcal{A}_{\varphi}}(q, x'_{k+1})$ and thus constraint (**Cau**) holds.

Theorem 2 thus holds for $\sigma \cdot (x_{k+1}, y_{k+1})$ in this case.

Thus Theorem 2 holds for $\sigma \cdot (x_{k+1}, y_{k+1})$.

B Appendix: Transformation of Non-Enforceable Properties

Let us recall the discussion about non-enforceable properties in Section 3 (Example 3, and the condition for enforceability (**EnfCo**)). We also saw that some non-enforceable properties can be transformed in to enforceable properties (by excluding some behaviors from the given non-enforceable property) via an example discussed in Remark 5, and an algorithm for transformation of non-enforceable properties is also briefly discussed in Section 3 after Remark 5.

Let us now discuss in detail about an algorithm that takes a safety automaton \mathcal{A}_{φ} that does not satisfy the condition for enforceability (**EnfCo**) and checks whether \mathcal{A}_{φ} can be transformed into an enforceable property (by excluding some behaviors) or not. If \mathcal{A}_{φ} can be transformed, then the algorithm returns the transformed safety automaton $\text{sub}(\mathcal{A}_{\varphi})$. The algorithm excludes only problematic paths (behaviors) from \mathcal{A}_{φ} , and all good behaviors will be retained in $\text{sub}(\mathcal{A}_{\varphi})$ (i.e., removal of behaviors is done minimally).

Algorithm 2 TransformNonEnf

```
1:  $\text{sub}(\mathcal{A}_\varphi) \leftarrow \mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$ 
2: while  $\exists q \in Q \setminus \{q_v\} : \forall (x, y) \in \Sigma, q \xrightarrow{(x,y)} q_v$  do
3:   for all  $q \in Q \setminus \{q_v\}$  do
4:     if  $\forall (x, y) \in \Sigma, q \xrightarrow{(x,y)} q_v$  then
5:       if  $q = q_0$  then
6:         RETURN(NONE)
7:       else
8:         remove( $q$ )
9:       end if
10:    end if
11:  end for
12: end while
13: RETURN( $\text{sub}(\mathcal{A}_\varphi)$ )
```

Algorithm 2 takes an SA $\mathcal{A}_\varphi = (Q, q_0, q_v, \Sigma, \rightarrow)$ (that does not satisfy the condition for enforceability) as input and returns an SA $\text{sub}(\mathcal{A}_\varphi)$ which is enforceable or it returns NONE if \mathcal{A}_φ cannot be transformed into an enforceable property. $\text{sub}(\mathcal{A}_\varphi)$ is initialized with the \mathcal{A}_φ . Function remove takes a state $q \in Q \setminus \{q_v\}$ and merges it with q_v (i.e., q is removed from the set of states Q and all the incoming transitions to q go to q_v instead).

The algorithm proceeds as follows: The condition of the while loop tests whether there are any states in $\text{sub}(\mathcal{A}_\varphi)$ that have all its outgoing transitions to q_v . If this condition evaluates to true, then each state q in $Q \setminus \{q_v\}$ is checked (whether all the outgoing transitions from q go to q_v). If $q = q_0$ is such a state (i.e., all the outgoing transitions from q_0 go to q_v), then the algorithm immediately returns “NONE” (i.e., that \mathcal{A}_φ cannot be transformed into an enforceable property). Otherwise, if q is different from q_0 and if all the outgoing transitions from q go to q_v , then state q is removed and merged with q_v . Finally, when there are no states in $Q \setminus \{q_v\}$ with all outgoing transitions to q_v , the while loop ends and the transformed automaton $\text{sub}(\mathcal{A}_\varphi)$ is returned.

Note that if the automaton that is given as input to the algorithm already satisfies the condition for enforceability (**EnfCo**), then the algorithm returns the same automaton (the while loop condition test will evaluate to false and thus is never executed).

Example 4. Let us now consider some examples. Consider the example non-enforceable property discussed in Section 3, presented in Figure 3. If the property automaton in Figure 3 is given as input to Algorithm 2, in the first iteration, the while condition test will evaluate to true since there is a state q_1 , such that all the outgoing transitions from q_1 go to q_v (i.e., $q_1 \xrightarrow{\Sigma} q_v$). Thus, in the first iteration, q_1 will be merged with q_v (i.e., q_1 will be removed and all the incoming transitions to q_1 go to q_v). Before the second iteration of the while loop starts, we will have only two locations q_0 and q_v in the automaton, where all the transitions from location q_0 go to q_v . In the second iteration the while condition will evaluate to true, and the algorithm returns NONE, since $q_0 \xrightarrow{\Sigma} q_v$ and the initial locations also needs to be removed.

Example 5. We already discussed in Section 3 that the example non-enforceable presented in Figure 4a can be transformed to an enforceable property. The automaton in Figure 4b presents the transformed automaton returned by Algorithm 2.